

Distributed Learning for Cooperative Inference

Angelia Nedić · Alex Olshevsky · César A. Uribe[†]

the date of receipt and acceptance should be inserted later

Abstract We study the problem of cooperative inference where a group of agents interact over a network and seek to estimate a joint parameter that best explains a set of observations. Agents do not know the network topology or the observations of other agents. We explore a variational interpretation of the Bayesian posterior density, and its relation to the stochastic mirror descent algorithm, to propose a new distributed learning algorithm. We show that, under appropriate assumptions, the beliefs generated by the proposed algorithm concentrate around the true parameter exponentially fast. We provide explicit non-asymptotic bounds for the convergence rate. Moreover, we develop explicit and computationally efficient algorithms for observation models belonging to exponential families.

1 Introduction

The increasing amount of data generated by recent applications of distributed systems such as social media, sensor networks, and cloud-based databases has brought considerable attention to distributed data processing approaches, in particular the design of distributed algorithms that take into account the communication constraints and make coordinated decisions in a distributed manner [27, 55, 3, 47, 4, 9, 63, 18, 10, 14, 22]. In a distributed system, the interactions between agents are usually restricted to follow certain constraints on the flow of information imposed by the network structure. Such information constraints cause the agents to only be able to use locally available information. This contrasts with centralized approaches where all information and computation resources are available at a single location [24, 68, 64, 62].

One traditional problem in decision-making is that of parameter estimation or statistical learning. Given a set of noisy observations coming from a joint distribution one would like to estimate a parameter or distribution that minimizes a certain loss function. For example, Maximum a Posteriori (MAP) or Minimum Least Squared Error (MLSE) estimators fit a parameter to some model of the observations. Both, MAP and MLSE estimators require some form of Bayesian posterior computation based on models that explain the observations for a given parameter. Computation of such a posteriori distributions depends on having exact models about the likelihood of the corresponding observations. This is one of the main difficulties of using

[†]This research is supported partially by the National Science Foundation under grant no. CPS 15-44953 and by the Office of Naval Research under grant no. N00014-17-1-2195. A preliminary version of some results presented in this paper were published in the IEEE Conference of Decision and Control 2016 [44].

A. Nedić
ECE Department, Arizona State University
E-mail: angelia.nedich@asu.edu

A. Olshevsky
Department of ECE and Division of Systems Engineering, Boston University
E-mail: alexols@bu.edu

C.A. Uribe ✉
ECE Department and Coordinated Science Laboratory, University of Illinois
E-mail: cauribe2@illinois.edu

Bayesian approaches in a distributed setting. A fully Bayesian approach is not possible because full knowledge of the network structure, or of other agents' likelihood models, may not be available [16, 37, 2].

Following the seminal work of Jadbabaie et al. in [27, 28, 58], there have been many studies of distributed non-Bayesian update rules over networks. In this case, agents are assumed to be boundedly rational (i.e., they fail to aggregate information in a fully Bayesian way [23]). Proposed non-Bayesian algorithms involve an aggregation step, typically consisting of weighted geometric or arithmetic average of the received beliefs [1, 63, 26, 40, 48], and a Bayesian update with the locally available data [2, 38]. Recent studies proposed variations of the non-Bayesian approach and proved consistent, geometric and non-asymptotic convergence rates for a general class of distributed algorithms; from asymptotic analysis [58, 31, 49, 50, 59, 54] to non-asymptotic bounds [60, 41, 32, 42], time-varying directed graphs [45], and transmission and node failures [61]; see [5, 46] for an extended literature review.

We build upon the work in [8] on non-asymptotic behaviors of Bayesian estimators to derive new non-asymptotic concentration results for distributed learning algorithms. In contrast to the existing results which assume a finite hypothesis set, in this paper we extend the framework to countably many and a continuum of hypotheses. Our results show that in general, the network structure will induce a transient time after which all agents learn at a network independent rate, and this rate is geometric.

The contributions of this paper are as follows. We begin with a variational analysis of Bayesian posterior and derive an optimization problem for which the posterior is a step of the Stochastic Mirror Descent method. We then use this interpretation to propose a distributed Stochastic Mirror Descent method for distributed learning. We show that this distributed learning algorithm concentrates the beliefs of all agents around the true parameter at an exponential rate. We derive high probability non-asymptotic bounds for the convergence rate. In contrast to the existing literature, we analyze the case where the parameter spaces are compact. Moreover, we specialize the proposed algorithm to parametric models of an exponential family which results in especially simple updates.

The rest of this paper is organized as follows. Section 2 introduces the problem setup, it describes the networked observation model and the inference task. Section 3 presents a variational analysis of the Bayesian posterior, shows the implicit representation of the posterior as steps in a stochastic program and extends this program to the distributed setup. Section 4 specializes the proposed distributed learning protocol to the case of observation models that are members of the exponential family. Section 5 shows our main results about the exponential concentration of beliefs around the true parameter. Section 5 begins by gently introducing our techniques by proving a concentration result in the case of countably many hypotheses, before turning to our main focus: the case when the set of hypotheses is a compact subset of \mathbb{R}^d . Finally, conclusions, open problems, and potential future work are discussed.

Notation: Random variables are denoted with upper-case letters, e.g. X , while the corresponding lower-case are used for their realizations, e.g. x . Time indices are denoted by subscripts, and the letter k or t is generally used. Agent indices are denoted by superscripts, and the letters i or j are used. We write $[A]_{ij}$ or a_{ij} to denote the entry of a matrix A in its i -th row and j -th column. We use A' for the transpose of a matrix A , and x' for the transpose of a vector x . The complement of a set B is denoted as B^c .

2 Problem Setup

We begin by introducing the learning problem from a centralized perspective, where all information is available at a single location. Later, we will generalize the setup to the distributed setting where only partial and distributed information is available.

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a sample space, \mathcal{F} is a σ -algebra and \mathbb{P} a probability measure. Assume that we observe a sequence of independent random variables X_1, X_2, \dots , all taking values in some measurable space $(\mathcal{X}, \mathcal{A})$ and identically distributed with a common *unknown* distribution P . In addition, we have a parametrized family of distributions $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$, where the map $\Theta \rightarrow \mathcal{P}$ from parameter to distribution is one-to-one. Moreover, the models in \mathcal{P} are all dominated¹ by a σ -finite measure λ , with corresponding densities $p_\theta = dP_\theta/d\lambda$. Assuming that there exists a θ^* such that $P_{\theta^*} = P$, the objective is to estimate θ^* based on the received observations x_1, x_2, \dots .

Following a Bayesian approach, we begin with a prior on θ^* represented as a distribution on the space Θ ; then given a sequence of observations, we incorporate such knowledge into a posterior distribution following Bayes' rule. Specifically, we assume that Θ is equipped with a σ -algebra and a measure σ and that μ_0 , which is our prior belief, is a probability measure on Θ which is dominated by σ . Furthermore, the densities $p_\theta(x)$ are measurable functions of θ for any $x \in \mathcal{X}$, and also dominated

¹ A measure μ is dominated by (or absolutely continuous with respect to) a measure λ if $\lambda(B) = 0$ implies $\mu(B) = 0$ for every measurable set B .

by σ . We then define the belief μ_k as the posterior distribution given the sequence of observations up to time k , i.e.,

$$\mu_{k+1}(B) \propto \int_B \prod_{t=1}^{k+1} p_\theta(x_t) d\mu_0(\theta). \quad (1)$$

for any measurable set $B \subset \Theta$ (note that we used the independence of the observations at each time step). Assuming that all observations are readily available at a centralized location, under appropriate conditions, the recursive Bayesian posterior in Eq. (1) will be consistent in the sense that the beliefs μ_k will concentrate around θ^* ; see [19, 57, 20] for a formal statement. Several authors have studied the rate at which this concentration occurs, in both asymptotic and non-asymptotic regimes [8, 21, 56].

Now consider the case where there is a network of n agents observing the process X_1, X_2, \dots , where X_k is now a random vector belonging to the product space $\prod_{i=1}^n \mathcal{X}^i$, and $X_k = [X_k^1, X_k^2, \dots, X_k^n]'$ consists of observations X_k^i of the agents at time k . Specifically, agent i observes the sequence X_1^i, X_2^i, \dots , where X_k^i is now distributed according to an unknown distributions P^i . Each agent i has a private family of distributions $\mathcal{P}^i = \{P_\theta^i : \theta \in \Theta\}$ it would like to fit to the observations. However, the goal is for *all* agents to agree on a *single* θ that best explains the complete set of observations. In other words, the agents collaboratively seek to find a θ^* that makes the distribution $P_{\theta^*} = \prod_{i=1}^n P_{\theta^*}^i$ as close as possible to the unknown true distribution $P = \prod_{i=1}^n P^i$. Agents interact over a network defined by an undirected graph $\mathcal{G} = (V, E)$, where $V = \{1, 2, \dots, n\}$ is the set of agents and E is a set of undirected edges, i.e., $(i, j) \in E$ if and only if agents i and j can communicate with each other.

We study a simple interaction model where, at each step, agents exchange their beliefs with their neighbors in the graph. Thus at every time step k , agent i will receive the sample x_k^i from X_k^i as well as the beliefs of its neighboring agents, i.e., it will receive μ_{k-1}^j for all j such that $(i, j) \in E$. Applying a fully Bayesian approach runs into some obstacles in this setting, as agents know neither the network topology nor the private family of distributions of other agents. Our goal is to design a learning procedure which is both distributed and consistent. That is, we are interested in a belief update algorithm that aggregates information in a non-Bayesian manner and guarantees that the beliefs of all agents will concentrate around θ^* .

As a motivating example, consider the problem of distributed source localization [52, 53]. In this scenario, a network of n agents receives noisy measurements of the distance to a source. The sensing capabilities of each sensor might be limited to a certain region. The group objective is to jointly identify the location of the source. Figure 1 shows a group of 7 agents (circles) seeking to localize a source (star). There is an underlying graph that indicates which nodes can exchange messages. Moreover, each node has a sensing region indicated by the dashed circle around it. Each agent observes signals proportional to the distance to the target. Since a target cannot be localized effectively from a single measure of the distance, agents must cooperate to have any hope of achieving decent localization. For more details on the problem, as well as simulations of the several discrete learning rules, we refer the reader to our earlier paper [41] dealing with the case when the set Θ is finite.




Fig. 1: Distributed source localization example.

3 A variational approach to distributed Bayesian filtering

In this section, we make the observation that the posterior in Eq. (1) corresponds to an iteration of a first-order optimization algorithm, namely Stochastic Mirror Descent [7, 39, 11, 51]. Closely related variational interpretations of Bayes' rule are well-known, and in particular have been given in [67, 65, 25]. The specific connection to Stochastic Mirror Descent has not been noted, as far as we are aware of. This connection will serve to motivate a distributed learning method which will be the main focus of the paper.

3.1 Bayes' rule as Stochastic Mirror Descent

Suppose we want to solve the following optimization problem

$$\min_{\theta \in \Theta} F(\theta) = D_{KL}(P \| P_\theta), \quad (2)$$

where P is an unknown true distribution and P_θ is a parametrized family of distributions (see Section 2). Here, $D_{KL}(P \| Q)$ is the Kullback-Leibler (KL) divergence² between distributions P and Q .

First note that we can rewrite Eq. (2) as

$$\begin{aligned} \min_{\theta \in \Theta} D_{KL}(P \| P_\theta) &= \min_{\pi \in \Delta_\Theta} \mathbb{E}_\pi D_{KL}(P \| P_\theta) \quad \text{s.t. } \theta \sim \pi \\ &= \min_{\pi \in \Delta_\Theta} \mathbb{E}_\pi \mathbb{E}_P \left[-\log \frac{dP_\theta}{dP} \right], \end{aligned}$$

where Δ_Θ is the set of all possible densities on the parameter space Θ . Since the distribution P does not depend on the parameter θ , it follows that

$$\begin{aligned} \arg \min_{\theta \in \Theta} D_{KL}(P \| P_\theta) &= \arg \min_{\pi \in \Delta_\Theta} \mathbb{E}_\pi \mathbb{E}_P [-\log p_\theta(X)] \quad \text{where } \theta \sim \pi \text{ and } X \sim P \\ &= \arg \min_{\pi \in \Delta_\Theta} \mathbb{E}_P \mathbb{E}_\pi [-\log p_\theta(X)] \quad \text{where } \theta \sim \pi \text{ and } X \sim P. \end{aligned} \quad (3)$$

The equality in Eq. (3), where we exchange the order of the expectations, follows from the Fubini-Tonelli theorem. Clearly, if θ^* minimizes Eq. (2), then a distributions which puts all the mass on θ^* minimizes Eq. (3).

The difficulty in evaluating the objective function in Eq. (3) lies in the fact that the distribution P is unknown. A generic approach to solving such problems is using algorithms from stochastic approximation methods, where the objective is minimized by constructing a sequence of gradient-based iterates whereby the true gradient of the objective (which is not available) is replaced with a gradient sample that is available at a given time.

A particular method that is relevant for the solution of stochastic programs of the form

$$\min_{x \in Z} \mathbb{E} [F(x, \Xi)],$$

for some random variable Ξ with unknown distribution, is the *stochastic mirror descent* method [29, 39, 7, 33]. The stochastic mirror descent approach constructs a sequence $\{x_k\}$ as follows:

$$x_{k+1} = \arg \min_{x \in Z} \left\{ \langle \nabla F(x, \xi_k), x \rangle + \frac{1}{\alpha_k} D_w(x, x_k) \right\},$$

for a realization ξ_k of Ξ . Here, $\alpha_k > 0$ is the step-size, $\langle p, q \rangle = \int_\Theta p(\theta)q(\theta)d\sigma$, and $D_w(x, x_k)$ is a Bregman distance function associated with a distance-generating function w , i.e.,

$$D_w(x, z) = w(z) - w(x) - \delta w[z; x - z],$$

where $\delta w[z; x - z]$ is the Fréchet derivative of w at z in the direction of $x - z$.

² $D_{KL}(P \| Q)$ between distributions P and Q (with P dominated by Q) is defined to be $D_{KL}(P \| Q) = -\mathbb{E}_P [\log dQ/dP]$.

For Eq. (3), Stochastic Mirror Descent generates a sequence of densities $\{d\mu_k\}$, as follows:

$$d\mu_{k+1} = \arg \min_{\pi \in \Delta_\Theta} \left\{ \langle -\log p_\theta(x_{k+1}), \pi \rangle + \frac{1}{\alpha_k} D_w(\pi, d\mu_k) \right\}, \quad \text{where } \theta \sim \pi. \quad (4)$$

If we choose $w(x) = \int x \log x$ as the distance-generating function, then the corresponding Bregman distance is the Kullback-Leibler (KL) divergence D_{KL} . Additionally, by selecting $\alpha_k = 1$, the solution to the optimization problem in Eq. (4) can be computed explicitly, where for each $\theta \in \Theta$,

$$d\mu_{k+1}(\theta) \propto p_\theta(x_{k+1})d\mu_k(\theta),$$

which is the particular definition for the posterior distribution according to Eq. (1) (a formal proof of this assertion is a special case of Proposition 1 shown later in the paper).

3.2 Distributed Stochastic Mirror Descent

Now, consider the distributed problem where the network of agents want to collectively solve the following optimization problem

$$\min_{\theta \in \Theta} F(\theta) \triangleq D_{KL}(\mathbf{P} \parallel \mathbf{P}_\theta) = \sum_{i=1}^n D_{KL}(P^i \parallel P_\theta^i). \quad (5)$$

Recall that the distribution \mathbf{P} is unknown (though, of course, agents gain information about it by observing samples from X_1^i, X_2^i, \dots and interacting with other agents) and that \mathcal{P}^i containing all the distributions P_θ^i is a private family of distributions and is only available to agent i .

We propose the following algorithm as a distributed version of the stochastic mirror descent for the solution of problem Eq. (5):

$$d\mu_{k+1}^i = \arg \min_{\pi \in \Delta_\Theta} \left\{ \langle -\log p_\theta^i(x_{k+1}^i), \pi \rangle + \sum_{j=1}^n a_{ij} D_{KL}(\pi \parallel d\mu_k^j) \right\} \quad \text{where } \theta \sim \pi, \quad (6)$$

with $a_{ij} > 0$ denoting the weight that agent i assigns to beliefs coming from its neighbor j . Specifically, $a_{ij} > 0$ if $(i, j) \in E$ or $j = i$, and $a_{ij} = 0$ if $(i, j) \notin E$. The optimization problem in Eq. (5) has a closed form solution. In particular, the posterior density at each $\theta \in \Theta$ is given by

$$d\mu_{k+1}^i(\theta) \propto p_\theta^i(x_{k+1}^i) \prod_{j=1}^n (d\mu_k^j(\theta))^{a_{ij}},$$

or equivalently, the belief on a measurable set B of an agent i at time $k+1$ is

$$\mu_{k+1}^i(B) \propto \int_B p_\theta^i(x_{k+1}^i) \prod_{j=1}^n (d\mu_k^j(\theta))^{a_{ij}}. \quad (7)$$

We state the correctness of this claim in the following proposition.

Proposition 1 *The probability measure μ_{k+1}^i over the set Θ defined by the update protocol Eq. (7) coincides, almost everywhere, with the update the distributed stochastic mirror descent algorithm applied to the optimization problem in Eq. (5).*

Proof We need to show that the density $d\mu_{k+1}^i$ associated with the probability measure μ_{k+1}^i defined by Eq. (7) minimizes the problem in Eq. (6). To do so, let $G(\pi)$ be the objective function for the problem in Eq. (6), i.e.,

$$G(\pi) = \langle -\log p_\theta^i(x_{k+1}^i), \pi \rangle + \sum_{j=1}^n a_{ij} D_{KL}(\pi \parallel d\mu_k^j).$$

Next, we add and subtract the KL divergence between π and the density $d\mu_{k+1}^i$ to obtain

$$\begin{aligned} G(\pi) &= \langle -\log p_\theta^i(x_{k+1}^i), \pi \rangle + \sum_{j=1}^n a_{ij} D_{KL}(\pi \| d\mu_k^j) - D_{KL}(\pi \| d\mu_{k+1}^i) + D_{KL}(\pi \| d\mu_{k+1}^i) \\ &= \langle -\log p_\theta^i(x_{k+1}^i), \pi \rangle + D_{KL}(\pi \| d\mu_{k+1}^i) + \sum_{j=1}^n a_{ij} \mathbb{E}_\pi \log \frac{d\mu_{k+1}^i}{d\mu_k^j}. \end{aligned}$$

Now, from Eq. (7) it follows that

$$\begin{aligned} G(\pi) &= \langle -\log p_\theta^i(x_{k+1}^i), \pi \rangle + D_{KL}(\pi \| d\mu_{k+1}^i) + \\ &\quad \sum_{j=1}^n a_{ij} \mathbb{E}_\pi \log \left(\frac{1}{d\mu_k^j} \frac{1}{Z_{k+1}^i} \prod_{l=1}^n (d\mu_k^l)^{a_{il}} p_\theta^i(x_{k+1}^i) \right) \\ &= \langle -\log p_\theta^i(x_{k+1}^i), \pi \rangle + D_{KL}(\pi \| d\mu_{k+1}^i) \\ &\quad - \log Z_{k+1}^i + \langle \log p_\theta^i(x_{k+1}^i), \pi \rangle + \sum_{j=1}^n a_{ij} \mathbb{E}_\pi \log \left(\frac{1}{d\mu_k^j} \prod_{l=1}^n (d\mu_k^l)^{a_{il}} \right) \\ &= -\log Z_{k+1}^i + D_{KL}(\pi \| d\mu_{k+1}^i) - \sum_{j=1}^n a_{ij} \mathbb{E}_\pi \log d\mu_k^j + \sum_{l=1}^n a_{il} \mathbb{E}_\pi \log d\mu_k^l \\ &= -\log Z_{k+1}^i + D_{KL}(\pi \| d\mu_{k+1}^i) \end{aligned} \tag{8}$$

where $Z_{k+1}^i = \int_\theta p_\theta^i(x_{k+1}^i) \prod_{j=1}^n (d\mu_k^j(\theta))^{a_{ij}}$ is the corresponding normalizing constant.

The first term in Eq. (8) does not depend on the distribution π . Thus, we conclude that the solution to the problem in Eq. (6) is the density $\pi^* = d\mu_{k+1}^i$ as defined in Eq. (7) (almost everywhere). \square

We remark that the update in Eq. (7) can be viewed as two-step processes: first every agent constructs an aggregate belief using a weighted geometric average of its own belief and the beliefs of its neighbors, and then each agent performs a Bayes' update using the aggregated belief as a prior. We note that similar arguments in the context of distributed optimization have been proposed in [51, 36] for general Bregman distances. In the case when the number of hypotheses is finite, variations on this update rule were previously analyzed in [60, 41, 32].

3.3 An example

Example 1 Consider a group of 4 agents, connected over a network as shown in Figure 2. A set of metropolis weights for this network is given by the following matrix:

$$A = \begin{bmatrix} 2/3 & 1/6 & 0 & 1/6 \\ 1/6 & 2/3 & 1/6 & 0 \\ 0 & 1/6 & 2/3 & 1/6 \\ 1/6 & 0 & 1/6 & 2/3 \end{bmatrix}.$$

Furthermore, assume that each agent is observing a Bernoulli random variable such that $X_k^1 \sim \text{Bern}(0.2)$, $X_k^2 \sim \text{Bern}(0.4)$, $X_k^3 \sim \text{Bern}(0.6)$ and $X_k^4 \sim \text{Bern}(0.8)$. In this case, the parameter space is $\Theta = [0, 1]$. Thus, the objective is to collectively find a parameter θ^* that best explains the joint observations in the sense of the problem in Eq. (5), i.e.

$$\min_{\theta \in [0, 1]} F(\theta) = \sum_{j=1}^4 D_{KL}(\text{Bern}(\theta) \| \text{Bern}(\theta^j)) = \sum_{j=1}^4 \left(\theta \log \frac{\theta}{\theta^j} + (1 - \theta) \log \frac{1 - \theta}{1 - \theta^j} \right)$$

where $\theta^1 = 0.2$, $\theta^2 = 0.4$, $\theta^3 = 0.6$ and $\theta^4 = 0.8$. We can see that the optimal solution is $\theta^* = 0.5$ by determining it explicitly via the first-order optimality conditions or by exploiting the symmetry in the objective function.




Fig. 2: A network of 4 agents.

Assume that all agents start with a common belief at time 0 following a Beta distribution, i.e., $\mu_0^i = \text{Beta}(\alpha_0, \beta_0)$ (this specific choice will be motivated in the next section). Then, the proposed algorithm in Eq. (7) will generate a belief at time $k + 1$ that also has a Beta distribution. Moreover, $\mu_{k+1}^i = \text{Beta}(\alpha_{k+1}^i, \beta_{k+1}^i)$, where

$$\alpha_{k+1}^i = \sum_{j=1}^n a_{ij} \alpha_k^j + x_{k+1}^i, \quad \beta_{k+1}^i = \sum_{j=1}^n a_{ij} \beta_k^j + 1 - x_{k+1}^i.$$

To summarize, we have given an interpretation of Bayes' rule as an instance of Stochastic Mirror Descent. We have shown how this interpretation motivates a distributed update rule. In the next section, we discuss explicit forms of this update rule for parametric models coming from exponential families.

4 Cooperative Inference for Exponential Families

We begin with the observation that, for a general class of models $\{\mathcal{P}^i\}$, it is not clear whether the computation of the posterior beliefs μ_{k+1}^i is tractable. Indeed, computation of μ_{k+1}^i involves solving an integral of the form

$$\int_{\Theta} p_{\theta}^i(x_{k+1}^i) \prod_{j=1}^n (d\mu_k^j(\theta))^{a_{ij}}. \quad (9)$$

There is an entire area of research called *variational Bayes' approximations* dedicated to efficiently approximating integrals that appear in such context [15, 6, 12].

The purpose of this section is to show that for exponential family [30, 13] there are closed-form expressions for the posteriors.

Definition 1 The exponential family, for a parameter $\theta = [\theta^1, \theta^2, \dots, \theta^s]'$, is the set of probability distributions whose density can be represented as

$$p_{\theta}(x) = H(x) \exp(M(\theta)'T(x) - C(\theta))$$

for specific functions $H(\cdot)$, $M(\cdot)$, $T(\cdot)$ and $C(\cdot)$, with $M(\theta) = [M(\theta^1), M(\theta^2), \dots, M(\theta^s)]'$. The function $M(\theta)$ is usually referred to as the natural parameter.

When $M(\theta)$ is used as a parameter itself, it is said that the distribution is in its canonical form. In this case, we can write the density as

$$p_M(x) = H(x) \exp(M'T(x) - C(M)),$$

with M being the parameter.

Among the members of the exponential family, one can find the distributions such as Normal, Poisson, Exponential, Gamma, Bernoulli, and Beta, among others [17]. In our case, we will take advantage of the existence of *conjugate priors* for all members of the exponential family. The definition of the conjugate prior is given below.

Definition 2 Assume that the prior distribution p on a parameter space Θ belongs to the exponential family. Then, the distribution p is referred to as the *conjugate prior* for a likelihood function $p_{\theta}(x)$ if the posterior distribution $p(\theta|x) \propto p_{\theta}(x)p(\theta)$ is in the same family as the prior.

Thus, if the belief density at some time k is a conjugate prior for our likelihood model, then our belief at time $k + 1$ will be of the same class as our prior. For example, if a likelihood function follows a Gaussian form, then having a Gaussian prior will produce a Gaussian posterior. This property simplifies the structure of the belief update procedure, since we can express the evolution of the beliefs generated by the proposed algorithm in Eq. (7) by the evolution of the natural parameters of the member of the exponential family it belongs to.

We now proceed to provide more details. First, the conjugate prior for a member of the exponential family can be written as

$$p_{\chi, \nu}(M) = f(\chi, \nu) \exp(M' \chi - \nu C(M)),$$

which is a distribution over the natural parameters M , where $\nu > 0$ and $\chi \in \mathbb{R}^s$ are the parameters of the conjugate prior. Then, it can be shown that the posterior distribution, given some observation x , has the same exponential form as the prior with updated parameters as follows:

$$p_{\chi, \nu}(M|x) = p_{\chi+T(x), \nu+1}(M) \propto p_{\theta}(x) p_{\chi, \nu}(M|x). \quad (10)$$

On the other hand, for a set on n priors of the same exponential family, the weighted geometric averages also have a closed form in terms of the conjugate parameters.

Proposition 2 *Let $(p_{\chi^1, \nu^1}(M), \dots, p_{\chi^n, \nu^n}(M))$ be a set of n distributions, all in the same class in the exponential family, i.e., $p_{\chi^i, \nu^i}(M) = f(\chi^i, \nu^i) \exp(M' \chi^i - \nu^i C(M))$ for $i = 1, \dots, n$. Then, for a set $(\alpha_1, \dots, \alpha_n)$ of weights with $\alpha_i > 0$ for all i , the probability distribution defined as*

$$p_{\bar{\chi}, \bar{\nu}}(M) = \frac{\prod_{i=1}^n (p_{\chi^i, \nu^i}(M))^{\alpha_i}}{\int \prod_{j=1}^n (p_{\chi^j, \nu^j}(dM))^{\alpha_j}},$$

belongs to the same class in the exponential family with parameters $\bar{\chi} = \sum_{i=1}^n \alpha_i \chi^i$ and $\bar{\nu} = \sum_{i=1}^n \alpha_i \nu^i$.

Proof We write the explicit geometric product, and discard the constant terms

$$\begin{aligned} p_{\bar{\chi}, \bar{\nu}}(M) &\propto \prod_{i=1}^n (f(\chi^i, \nu^i) \exp(M' \chi^i - \nu^i C(M)))^{\alpha_i} \\ &\propto \exp \left(M' \sum_{i=1}^n \alpha_i \chi^i - \sum_{i=1}^n \alpha_i \nu^i C(M) \right). \end{aligned}$$

The last line provides explicit values for the parameters of the new distribution. \square

The relations in Eq. (10) and Proposition 2 allow us to write the algorithm in Eq. (7) in terms of the natural parameters of the priors, as shown by the following proposition.

Proposition 3 *Assume that the belief density $d\mu_k^i$ at time k has an exponential form with natural parameters χ_k^i and ν_k^i for all $1 \leq i \leq n$, and that these densities are conjugate priors of the likelihood models p_{θ}^i . Then, the belief density at time $k + 1$, as computed in the update rule in Eq. (7), has the same form as the beliefs at time k with the natural parameters given by*

$$\chi_{k+1}^i = \sum_{j=1}^n a_{ij} \chi_k^j + T^i(x^i), \quad \nu_{k+1}^i = \sum_{j=1}^n a_{ij} \nu_k^j + 1 \quad \text{for all } i = 1, \dots, n.$$

The proof of Proposition 3 follows immediately from Eq. (10) and Eq. (2).

Proposition 3 simplifies the algorithm in Eq. (7) and facilitates its use in traditional estimation problems where members of the exponential family are used. We next illustrate this by discussing a number of distributed estimation problems with likelihood models coming from exponential families.

4.1 Distributed Poisson Filter

Consider an observation model where the agent signals follow Poisson distributions, i.e., $X_k^i = \text{Poisson}(\lambda^i)$ for all i . In this case, the optimization problem to be solved is

$$\min_{\lambda > 0} F(\lambda) = \sum_{j=1}^n D_{KL}(\text{Poisson}(\lambda) \parallel \text{Poisson}(\lambda^j)),$$

or equivalently, $\min_{\lambda > 0} \{-\sum_{i=1}^n \lambda^i \log \lambda + \lambda\}$.

The conjugate prior of a Poisson likelihood model is the Gamma distribution. Thus, if at time k the beliefs are given by $\mu_k^i = \text{Gamma}(\alpha_k^i, \beta_k^i)$ for all i , then the beliefs at time $k+1$ are $\mu_{k+1}^i = \text{Gamma}(\alpha_{k+1}^i, \beta_{k+1}^i)$, where

$$\alpha_{k+1}^i = \sum_{j=1}^n a_{ij} \alpha_k^j + x_{k+1}^i \quad \text{and} \quad \beta_{k+1}^i = \sum_{j=1}^n a_{ij} \beta_k^j + 1.$$

4.2 Distributed Gaussian Filter with known variance

Assume each agent observes a signal of the form $X_k^i = \theta^i + \epsilon_k^i$, where θ^i is finite and unknown, while $\epsilon^i \sim \mathcal{N}(0, 1/\tau^i)$, with $\tau^i = 1/(\sigma^i)^2$, is known by agent i . The optimization problem to be solved is

$$\min_{\theta \in \mathbb{R}} F(\theta) = \sum_{j=1}^n D_{KL}(\mathcal{N}(\theta, 1/\tau^j) \parallel \mathcal{N}(\theta^j, 1/\tau^j)),$$

or equivalently $\min_{\theta \in \mathbb{R}} \sum_{j=1}^n \tau^j (\theta - \theta^j)^2$.

In this case, the likelihood models, the prior and the posterior are Gaussian. Thus, if the beliefs of the agents at time k are Gaussian, i.e., $\mu_k^i = \mathcal{N}(\theta_k^i, 1/\tau_k^i)$ for all $i = 1 \dots, n$, then their beliefs at time $k+1$ are also Gaussian. In particular, they are given by $\mu_{k+1}^i = \mathcal{N}(\theta_{k+1}^i, 1/\tau_{k+1}^i)$ for all $i = 1 \dots, n$, with

$$\tau_{k+1}^i = \sum_{j=1}^n a_{ij} \tau_k^j + \tau^i \quad \text{and} \quad \theta_{k+1}^i = \frac{1}{\tau_{k+1}^i} \left(\sum_{j=1}^n a_{ij} \tau_k^j \theta_k^j + x_{k+1}^i \tau^i \right).$$

We note that this specific setup is known as Gaussian Learning and has been studied in [43, 66], where the expected parameter estimator is shown to converge at an $O(1/k)$ rate.

4.3 Distributed Gaussian Filter with unknown variance

In this case, the agents want to cooperatively estimate the value of a variance. Specifically, based on observations of the form $X_k^i = \theta^i + \epsilon_k^i$, with $\epsilon_k^i \sim \mathcal{N}(0, 1/\tau^i)$, where θ^i is known and τ^i is unknown to agent i , they want to solve the following problem

$$\min_{\tau > 0} F(\tau) = \sum_{j=1}^n D_{KL}(\mathcal{N}(\theta^j, 1/\tau) \parallel \mathcal{N}(\theta^j, 1/\tau^j)).$$

We choose the Scaled Inverse Chi-Squared³ as the distribution of our prior, so that $\mu_k^i = \text{Scaled Inv-}\chi^2(\nu_k^i, \tau_k^i)$ for all i , then the beliefs at time $k+1$ are given by $\mu_{k+1}^i = \text{Scaled Inv-}\chi^2(\nu_{k+1}^i, \tau_{k+1}^i)$ for all i , with

$$\nu_{k+1}^i = \sum_{j=1}^n a_{ij} \nu_k^j + 1 \quad \text{and} \quad \tau_{k+1}^i = \frac{1}{\nu_{k+1}^i} \left(\sum_{j=1}^n a_{ij} \nu_k^j \tau_k^j + (x_{k+1}^i - \theta^i)^2 \right).$$

³ The density function of the Scaled Inverse Chi-Squared is defined for $x > 0$ as $p_{\nu, \tau}(x) = \frac{(\tau \nu / 2)^{\nu/2}}{\Gamma(\nu/2)} \frac{\exp(-\frac{\tau x}{2})}{x^{1+\nu/2}}$.

4.4 Distributed Gaussian Filter with unknown mean and variance

In the preceding examples, we have considered the cases when either the mean or the variance is known. Here, we will assume that both the mean and the variance are unknown and need to be estimated. Explicitly, we still have noise observations $X_k^i = \theta^i + \epsilon_k^i$, with $\epsilon_k^i \sim \mathcal{N}(0, 1/\tau^i)$, and want to solve

$$\min_{\theta \in \mathbb{R}, \tau > 0} F(\theta, \tau) = \sum_{j=1}^n D_{KL}(\mathcal{N}(\theta, 1/\tau) \parallel \mathcal{N}(\theta^j, 1/\tau^j)).$$

The Normal-Inverse-Gamma distribution serves as conjugate prior for the likelihood model over the parameters (θ, τ) . Specifically, we assume that the beliefs at time k are given by

$$\mu_k^i = \text{Normal-Inv-Gamma}(\theta_k^i, \tau_k^i, \alpha_k^i, \beta_k^i) \quad \text{for all } i = 1, \dots, n.$$

Then, the beliefs at time $k + 1$ will have a Normal-Inverse-Gamma distribution with the following parameters

$$\begin{aligned} \tau_{k+1}^i &= \sum_{j=1}^n a_{ij} \tau_k^j + 1, & \theta_{k+1}^i &= \frac{\sum_{j=1}^n a_{ij} \tau_k^j \theta_k^j + x_{k+1}^i}{\tau_{k+1}^i}, \\ \alpha_{k+1}^i &= \sum_{j=1}^n a_{ij} \alpha_k^j + 1/2, & \beta_{k+1}^i &= \sum_{j=1}^n a_{ij} \beta_k^j + \frac{\sum_{j=1}^n a_{ij} \tau_k^j (x_{k+1}^i - \theta_k^j)^2}{2\tau_{k+1}^i}. \end{aligned}$$

5 Belief Concentration Rates

We now turn to the presentation of our main results which concern the rate at which beliefs generated by the update rule in Eq. (7) concentrate around the true parameter θ^* . We will break up our analysis into two cases. Initially, we will focus on the case when Θ is a countable set, and will prove a concentration result for a ball containing the optimal hypothesis having finitely many hypotheses outside it. We will use this case to gently introduce the techniques we will use. We will then turn to our main scenario of interest, namely when Θ is a compact subset of \mathbb{R}^d . Our proof techniques use concentration arguments for beliefs on Hellinger balls from the recent work [8] which, in turn, builds on the classic paper [34].

We begin with two subsections focusing on background information, definitions, and assumptions.

5.1 Background: Hellinger Distance and Coverings

We equip the set of all probability distributions over the parameter set \mathcal{P} with the Hellinger distance⁴ to obtain the *metric* space (\mathcal{P}, h) . The metric space induces a topology, where we can define an open ball $\mathcal{B}_r(\theta)$ with a radius $r > 0$ centered at a point $\theta \in \Theta$, which we use to construct a special covering of subsets $B \subset \mathcal{P}$.

Definition 3 Define an n -Hellinger ball of radius r centered at θ as

$$\mathcal{B}_r(\theta) = \left\{ \hat{\theta} \in \Theta \mid \sqrt{\frac{1}{n} \sum_{i=1}^n h^2(P_{\theta}^i, P_{\hat{\theta}}^i)} \leq r \right\}.$$

Additionally, when no center is specified, it should be assumed that it refers to θ^* , i.e. $\mathcal{B}_r = \mathcal{B}_r(\theta^*)$.

⁴ The Hellinger distance between two probability distributions P and Q is given by,

$$h^2(P, Q) = \frac{1}{2} \int \left(\sqrt{\frac{dP}{d\lambda}} - \sqrt{\frac{dQ}{d\lambda}} \right)^2 d\lambda,$$

where P and Q are dominated by λ . Note that this formula is for the square of the Hellinger distance.

Given an n -Hellinger ball of radius r , we will use the following notation for a covering of its complement \mathcal{B}_r^c . Specifically, we are going to express \mathcal{B}_r^c as the union of finite disjoint and concentric annuli. Let $r > 0$ and $\{r_l\}$ be a finite strictly decreasing sequence such that $r_1 = 1$ and $r_L = r$. Now, express the set \mathcal{B}_r^c as the union of annuli generated by the sequence $\{r_l\}$ as

$$\mathcal{B}_r^c = \bigcup_{l=1}^{L-1} \mathcal{F}_l,$$

where $\mathcal{F}_l = \mathcal{B}_{r_l} \setminus \mathcal{B}_{r_{l+1}}$.

5.2 Background: Assumptions on Network and Mixing Weights

Naturally, we need some assumptions on the matrix A . For one thing, the matrix A has to be “compatible” with the underlying graph, in that information from node i should not affect node j if there is no edge from i to j in \mathcal{G} . At the other extreme, we want to rule out the possibility that A is the identity matrix, which in terms of Eq. (7) means nodes do not talk to their neighbors. Formally, we make the following assumption.

Assumption 1 *The graph \mathcal{G} and matrix A are such that:*

- (a) *A is doubly-stochastic with $[A]_{ij} = a_{ij} > 0$ for $i \neq j$ if and only if $(i, j) \in E$.*
- (b) *A has positive diagonal entries, $a_{ii} > 0$ for all $i \in V$.*
- (c) *The graph \mathcal{G} is connected.*

Assumption 1 is common in the distributed optimization literature. The construction of a set of weights satisfying Assumption 1 can be done in a distributed way, for example, by choosing the so-called “lazy Metropolis” matrix, which is a stochastic matrix given by

$$a_{ij} = \begin{cases} \frac{1}{2 \max\{d^i+1, d^j+1\}} & \text{if } (i, j) \in E, \\ 0 & \text{if } (i, j) \notin E, \end{cases}$$

where d^i is the degree (the number of neighbors) of node i . Note that although the above formula only gives the off-diagonal entries of A , it uniquely defines the entire matrix (the diagonal elements are uniquely defined via the stochasticity of A). To choose the weights corresponding to a lazy Metropolis matrix, agents will need to spend an additional round at the beginning of the algorithm broadcasting their degrees to their neighbors.

Assumption 1 can be seen to guarantee that $A^t \rightarrow (1/n)\mathbf{1}\mathbf{1}^T$ where $\mathbf{1}$ is the vector of all ones. We will use the following result that provides convergence rate for the difference $|A^t - (1/n)\mathbf{1}\mathbf{1}^T|$, based on the results from [60] and [41]:

Lemma 1 *Let Assumption 1 hold, then the matrix A satisfies the following relation:*

$$\sum_{t=1}^k \sum_{j=1}^n \left| \left[A^{k-t} \right]_{ij} - \frac{1}{n} \right| \leq \frac{4 \log n}{1 - \delta} \quad \text{for } i = 1, \dots, n,$$

where $\delta = 1 - \eta/4n^2$ with η being the smallest positive entry of the matrix A . Furthermore, if A is a lazy Metropolis matrix associated with the graph \mathcal{G} , then $\delta = 1 - 1/\mathcal{O}(n^2)$.

5.3 Concentration for the Case of Countable Hypotheses

We now turn to proving a concentration result when the set Θ of hypotheses is countable. We will consider the case of a ball in the Hellinger distance containing a countable number of hypotheses, including the correct one, and having only finitely many hypotheses outside it; we will show exponential convergence of beliefs to that ball. The purpose is to gently introduce the techniques we will use later in the case of a compact set of hypotheses.

In the case when the number of hypotheses is countable, the density update in Eq. (7) can be restated in a simpler form for discrete beliefs over the parameter space Θ as

$$\mu_{k+1}^i(\theta) \propto p_{\theta}^i(x_{k+1}^i) \prod_{j=1}^n (\mu_k^j(\theta))^{a_{ij}}. \quad (11)$$

We will fix the radius r , and our goal will be to prove a concentration result for a Hellinger ball of radius r around the optimal hypothesis θ^* . We partition the complement of this ball \mathcal{B}_r^c as described above into annuli \mathcal{F}_l . We introduce the notation \mathcal{N}_l to denote the number of hypotheses within the annulus \mathcal{F}_l . We refer the reader to Figure 3 which shows a set of probability distributions, represented as black dots, where the true distribution P is represented by a star.




Fig. 3: Creating a covering for a ball \mathcal{B}_r . \star represents the correct hypothesis P_{θ^*} , \bullet indicates the location of other hypotheses and the dash lines indicate the boundary of the balls \mathcal{B}_{r_l} .

We will assume that the number of hypotheses outside the desired ball is finite.

Assumption 2 *The number of hypothesis outside \mathcal{B}_r is finite.*

Additionally, we impose a bound on the separation between hypotheses which will avoid some pathological cases. The separation between hypotheses is defined in terms of the Hellinger affinity between two distributions Q and P , given by

$$\rho(Q, P) = 1 - h^2(Q, P).$$

Assumption 3 *There exists an $\alpha > 0$ such that $\rho(P_{\theta_1}^i, P_{\theta_2}^i) > \alpha$ for any $\theta_1, \theta_2 \in \Theta$ and $i = 1, \dots, n$.*

With these assumptions in place, our first step is a lemma that bounds concentration of log-likelihood ratios.

Lemma 2 *Let Assumptions 1, 2 and 3 hold. Given a set of independent random variables $\{X_t^i\}$ such that $X_t^i \sim P^i$ for $i = 1, \dots, n$ and $t = 1, \dots, k$, a set of distributions $\{Q^i\}$ where Q^i dominates P^i , then for all $y \in \mathbb{R}$,*

$$\mathbb{P} \left[\sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{dQ^j}{dP^j}(X_t^j) \geq y \right] \leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \exp \left(-k \frac{1}{n} \sum_{j=1}^n h^2(Q^j, P^j) \right).$$

Proof By the Markov inequality and Jensen's inequality we have

$$\begin{aligned} \mathbb{P} \left[\sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{dQ^j}{dP^j}(X_t^j) \geq y \right] &\leq \exp(-y/2) \mathbb{E} \left[\prod_{t=1}^k \prod_{j=1}^n \sqrt{\left(\frac{dQ^j}{dP^j}(X_t^j) \right)^{[A^{k-t}]_{ij}}} \right] \\ &\leq \exp(-y/2) \prod_{t=1}^k \prod_{j=1}^n \mathbb{E} \left[\sqrt{\left(\frac{dQ^j}{dP^j}(X_t^j) \right)^{[A^{k-t}]_{ij}}} \right] \end{aligned}$$

$$\leq \exp(-y/2) \prod_{t=1}^k \prod_{j=1}^n \rho(Q^j, P^j)^{[A^{k-t}]_{ij}},$$

where the last inequality follows from the definition of the Hellinger affinity function $\rho(Q, P)$. Now, by adding and subtracting $\frac{1}{n} \sum_{j=1}^n \log \rho(Q^j, P^j)$ we have

$$\begin{aligned} \mathbb{P} \left[\sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{dQ^j}{dP^j}(X_t^j) \geq y \right] &\leq \exp(-y/2) \exp \left(\sum_{t=1}^k \sum_{j=1}^n ([A^{k-t}]_{ij} - 1/n) \log \rho(Q^j, P^j) + \sum_{t=1}^k \frac{1}{n} \sum_{j=1}^n \log \rho(Q^j, P^j) \right) \\ &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \sum_{t=1}^k \sum_{j=1}^n |[A^{k-t}]_{ij} - 1/n| + \sum_{t=1}^k \frac{1}{n} \sum_{j=1}^n \log \rho(Q^j, P^j) \right), \end{aligned}$$

where the last line follows from $\rho(P^j, Q^j) > \alpha$.

Then, from Lemma 1 it follows that

$$\begin{aligned} \mathbb{P} \left[\sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{dQ^j}{dP^j}(X_t^j) \geq y \right] &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} + \sum_{t=1}^k \frac{1}{n} \sum_{j=1}^n \log \rho(Q^j, P^j) \right) \\ &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \prod_{j=1}^n \rho^k(Q^j, P^j)^{1/n} \\ &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \prod_{j=1}^n \exp(-kh^2(Q^j, P^j))^{1/n}. \end{aligned}$$

The last inequality follows from $\rho(Q^j, P^j) = 1 - h^2(Q^j, P^j)$ and $1 - x \leq \exp(-x)$ for $x \in [0, 1]$. \square

We are now ready to state our first main result, which bounds concentration of Eq. (11) around the optimal hypothesis for a countable hypothesis set Θ . The following theorem shows that the beliefs of all agents will concentrate around the Hellinger ball \mathcal{B}_r at an exponential rate.

Theorem 1 *Let Assumptions 1, 2 and 3 hold, and let $\sigma \in (0, 1)$ be a desired probability tolerance. Then, the belief sequences $\{\mu_k^i\}$, $i \in V$ that are generated by the update rule in Eq. (11), with initial beliefs such that $\mu_0^i(\theta^*) > \epsilon$ for all i , have the following property: for any radius $r > 0$ with probability $1 - \sigma$,*

$$\mu_{k+1}^i(\mathcal{B}_r) \geq 1 - \frac{1}{\epsilon} \chi \exp(-kr^2) \quad \text{for all } i \text{ and all } k \geq N,$$

where

$$N = \inf \left\{ t \geq 1 \mid \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \sum_{l=1}^{L-1} \mathcal{N}_{r_l} \exp(-tr_{l+1}^2) < \rho \right\},$$

$\chi = \sum_{l=1}^{L-1} \exp(-\frac{1}{2}r_l^2 + \log \mathcal{N}_{r_l})$, $\delta = 1 - \eta/n^2$, and η is the smallest positive element of the matrix A .

Proof We are going to focus on bounding the beliefs of a measurable set B , such that $\theta^* \in B$. For such a set, it follows from Eq. (11) that

$$\mu_k^i(B) = \frac{1}{Z_k^i} \sum_{\theta \in B} \left(\prod_{j=1}^n \mu_0^j(\theta)^{[A^k]_{ij}} \right) \prod_{t=1}^k \prod_{j=1}^n p_\theta^j(X_t^j)^{[A^{k-t}]_{ij}},$$

where Z_k^i is the appropriate normalization constant.

Furthermore, after a few algebraic operations we obtain

$$\mu_k^i(B) \geq 1 - \sum_{\theta \in B^c} \prod_{j=1}^n \left(\frac{\mu_0^j(\theta)}{\mu_0^j(\theta^*)} \right)^{[A^k]_{ij}} \prod_{t=1}^k \prod_{j=1}^n \left(\frac{p_\theta^j(X_t^j)}{p^j(X_t^j)} \right)^{[A^{k-t}]_{ij}}.$$

Moreover, since $\mu_0^i(\theta^*) > \epsilon$ for all $i = 1, \dots, n$, it follows that

$$\mu_k^i(B) \geq 1 - \frac{1}{\epsilon} \sum_{\theta \in B^c} \prod_{t=1}^k \prod_{j=1}^n \left(\frac{p_\theta^j(X_t^j)}{p^j(X_t^j)} \right)^{[A^{k-t}]_{ij}}. \quad (12)$$

The relation in Eq. (12) describes the iterative averaging of products of density functions, for which we can use Lemma 2 with $Q = P_\theta$ and $P = P_{\theta^*}$. Then,

$$\mathbb{P} \left(\left\{ \mathbf{X}^k \left| \sup_{\theta \in B^c} \sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{p_\theta^j(X_t^j)}{p^j(X_t^j)} \geq y \right. \right\} \right) \leq \sum_{\theta \in B^c} \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \exp \left(-k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right)$$

and by setting $y = -k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j)$ we obtain

$$\begin{aligned} & \mathbb{P} \left(\left\{ \mathbf{X}^k \left| \sup_{\theta \in B^c} \sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{p_\theta^j(X_t^j)}{p^j(X_t^j)} \geq -k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right. \right\} \right) \\ & \leq \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \sum_{\theta \in B^c} \exp \left(-\frac{k}{2} \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right). \end{aligned}$$

Now, we let the set B be the Hellinger ball of a radius r centered at θ^* and define a cover (as described above) to exploit the representation of B_r^c as the union of concentric Hellinger annuli, for which we have

$$\begin{aligned} & \mathbb{P} \left(\left\{ \mathbf{X}^k \left| \sup_{\theta \in B^c} \sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{p_\theta^j(X_t^j)}{p^j(X_t^j)} \geq -k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right. \right\} \right) \\ & \leq \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \sum_{l=1}^{L-1} \sum_{\theta \in \mathcal{F}_l} \exp \left(-\frac{k}{2} \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right) \\ & \leq \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \sum_{l=1}^{L-1} \mathcal{N}_{r_l} \exp \left(-kr_{l+1}^2 \right). \end{aligned}$$

We are interested in finding a value of k large enough such that the above probability is below σ . Thus, let's define the value of N as

$$N = \inf \left\{ t \geq 1 \left| \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \sum_{l=1}^{L-1} \mathcal{N}_{r_l} \exp \left(-tr_{l+1}^2 \right) < \sigma \right. \right\}.$$

It follows that for all $k \leq N$ with probability $1 - \sigma$, for all $\theta \in B_r^c$

$$\sum_{t=1}^k \sum_{j=1}^n [A^{k-t}]_{ij} \log \frac{p_\theta^j(X_t^j)}{p^j(X_t^j)} \leq -k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j).$$

Thus, from Eq. (12) with probability $1 - \sigma$ we have

$$\begin{aligned}
\mu_k^i(\mathcal{B}_r) &\geq 1 - \frac{1}{\epsilon} \sum_{\theta \in \mathcal{B}_r^c} \exp \left(-k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right) \\
&= 1 - \frac{1}{\epsilon} \sum_{l=1}^{L-1} \sum_{\theta \in \mathcal{F}_l} \exp \left(-k \frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j) \right) \\
&\geq 1 - \frac{1}{\epsilon} \sum_{l=1}^{L-1} \mathcal{N}_{r_l} \exp \left(-k r_{l+1}^2 \right) \\
&\geq 1 - \frac{1}{\epsilon} \sum_{l=1}^{L-1} \mathcal{N}_{r_l} \exp \left(-r_{l+1}^2 \right) \exp \left(-(k-1)r_l^2 \right) \\
&\geq 1 - \chi \frac{1}{\epsilon} \exp \left(-(k-1)r^2 \right),
\end{aligned}$$

where $\chi = \sum_{l=1}^{L-1} \exp \left(-\frac{1}{2}r_l^2 + \log \mathcal{N}_{r_l} \right)$. \square

5.4 A Concentration Result for a Compact Set of Hypotheses

Next we consider the case when the hypothesis set Θ is a compact subset of \mathbb{R}^d . We will now additionally require the map from Θ to $\prod_{i=1}^n P_\theta^i$ be continuous (where the topology on the space of distributions comes from the Hellinger metric). This will be useful in defining coverings, which will be made clear shortly.

Definition 4 Let (M, d) be a metric space. A subset $S \subseteq M$ is called ε -separated with $\varepsilon > 0$ if $d(x, y) \geq \varepsilon$ for any $x, y \in S$. Moreover, for a set $B \subseteq M$, let $N_B(\varepsilon)$ be the smallest number of Hellinger balls with centers in S of radius ε needed to cover the set B , i.e., such that $B \subseteq \bigcup_{m \in S} \mathcal{B}_\varepsilon(m)$.

As before, given a decreasing sequence $1 = r_1 \geq r_2 \geq \dots \geq r_L = r$, we will define the annulus \mathcal{F}_l to be $\mathcal{F}_l = \mathcal{B}_{r_l} \setminus \mathcal{B}_{r_{l+1}}$. Furthermore, S_{ε_l} will denote maximal ε_l -separated subset of \mathcal{F} . Finally, $K_l = |S_{\varepsilon_l}|$.

We note that, as a consequence of our assumption that the map from Θ to $\prod_{i=1}^n P_\theta^i$ is continuous, we have that each K_l is finite (since the image of a compact set under a continuous map is compact). Thus, we have the following covering of \mathcal{B}_r^c :

$$\mathcal{B}_r^c = \bigcup_{l=1}^{L-1} \bigcup_{m \in S_{\varepsilon_l}} \mathcal{F}_{l,m},$$

where each $\mathcal{F}_{l,m}$ is the intersection of a ball in S_{ε_l} with \mathcal{F}_l . Figure 4 shows the elements of a covering for a set \mathcal{B}_r^c . The cluster of circles at the top right corner represents the balls $\mathcal{B}_{\varepsilon_l}$ and, for a specific case in the left of the image, we illustrate the set $\mathcal{F}_{l,m}$.

Example 2 We continue Example 1 from Section 3. Suppose we are interested in analyzing the concentration of the beliefs around the true parameter θ^* on a Euclidean ball of radius 0.05; that is we want to see the total mass on the set $[0.45, 0.55]$. This in turn, represents a Hellinger ball of radius $r = 0.001254$. For this choice of r , we propose a covering where $r_1 = 1$, $r_2 = 1/2$, $r_3 = 1/4$, \dots , $r_{10} = 1/512$, $r_{11} = r$.

Figure 5 shows the Hellinger distance between the hypotheses p_θ and the optimal one p_{θ^*} . Specifically, the x -axis is the value of θ , and the y -axis shows the Hellinger distance between the distributions. Figure 5 also shows the covering we defined before, as horizontal lines for each value of the sequence r_l , which in turn defines the annulus \mathcal{F}_l . The Hellinger ball of radius r is also shown, with the corresponding subset of Θ where we want to analyze the belief concentration.

In this example, the parameter has dimension 1. The number of balls needed to cover each annulus can be seen to be 2, i.e., we only need 2 balls of radius $r_l/2$ to cover the annulus \mathcal{F}_l . Thus, $K_l = 2$ for $1 \leq l \leq L-1$. \square

Our concentration result requires the following assumption on the densities.




Fig. 4: Creating a covering for a set B_r . \star represents the correct hypothesis P_{θ^*} .




Fig. 5: Hellinger distance of the density p_θ to the optimal density p_{θ^*} .

Assumption 4 For every $i = 1, \dots, n$ and all θ , it holds that $p_\theta^i(x) \leq 1$ almost everywhere.

Assumption 4 will be technically convenient for us. It can be made without loss of generality in the following sense: we can always modify the underlying problem to make it hold.

Let us give an example before explaining the reasoning behind this assertion. Let us assume there is just one agent, and say $X \sim P$ is Gaussian with mean $\theta^* = 5$ and variance 0.01. Our model is $P_\theta = \mathcal{N}(\theta, 0.01)$ for $\theta \in \Theta = [0, 10]$. Because the variance is so small, the density values are larger than 1. Instead let us multiply all our observations by 10. We will then have that our observations come from $10X$, which indeed has density upper bounded by one. In turn our model now should be $Q_\theta = \mathcal{N}(10\theta, 1)$ or, alternatively, $Q_\theta = \mathcal{N}(\theta, 1)$ for $\theta \in \hat{\Theta} = [0, 100]$.

We note that this modification does not come without cost. As in the case of countable hypotheses, our convergence rates will depend on α , defined to be a positive number such that $\rho(P_{\theta^1}, P_{\theta^2}) > \alpha$ for any θ^1 and θ^2 . The process we have sketched out can decrease this parameter α .

In the general case, if each agent observe $X_t^j \sim P^j$, then there exists a large enough constant $M > 1$ such that $MX_t^j \sim Q^j$ where the density of Q^j is at most 1. We can then have agents multiply their measurements by M and redefine the densities to account for this.

We next provide a concentration result for the logarithmic likelihood of a ratio of densities, which will serve the same technical function as Lemma 2 in the countable hypothesis case. We begin by defining two measures. For a hypothesis θ and a measurable set $B \subseteq \Theta$, let $P_B^{\otimes k}$ be the probability distribution with density

$$g_B(\mathbf{x}^k) = \frac{1}{\mu_0(B)} \int_B \prod_{t=1}^k \prod_{j=1}^n p_\theta^j(x_t^j) d\mu_0(\theta). \quad (13)$$

Similarly, let $\bar{P}_B^{\otimes k}$ be the measure with density (i.e., Radon-Nikodym derivative with respect to $\lambda^{\otimes nk}$),

$$\bar{g}_B(\mathbf{x}^k) = \frac{1}{\mu_0(B)} \int_B \prod_{t=1}^k \prod_{j=1}^n (p_\theta^j(x_t^j))^{[A^{k-t}]_{ij}} d\mu_0(\theta). \quad (14)$$

Note that $\bar{P}_B^{\otimes k}$'s are not probability distributions due to the exponential weights. Nonetheless, they are bounded and positive. The next lemma shows the concentration of the logarithmic ratio of two weighted densities, as defined in Eq. (14), for two different sets B_1 and B_2 , in terms of the probability distribution $P_{B_1}^{\otimes k}$.

Lemma 3 *Let Assumptions 1, 3 and 4 hold. Consider two measurable sets $B_1, B_2 \subset \Theta$, both with positive measures, and assume that $B_1 \subset \mathcal{B}_{r^1}(\theta^1)$ and $B_2 \subset \mathcal{B}_{r^2}(\theta^2)$ where $\mathcal{B}_{r^1}(\theta^1)$ and $\mathcal{B}_{r^2}(\theta^2)$ are disjoint. Then, for all $y \in \mathbb{R}$*

$$\mathbb{P}_{B_1} \left[\log \frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)} \geq y \right] \leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \exp \left(-k \left(\sqrt{\frac{1}{n} \sum_{j=1}^n h^2(P_{\theta^1}^j, P_{\theta^2}^j)} - r^1 - r^2 \right)^2 \right),$$

where \mathbb{P}_{B_1} is the probability measure that gives \mathbf{X}^k having a distribution $P_{B_1}^{\otimes k}$ with density g_{B_1} as defined in Eq. (13).

Proof By the Markov inequality, it follows that

$$\begin{aligned} \mathbb{P}_{B_1} \left[\log \frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)} \geq y \right] &\leq \exp(-y/2) \mathbb{E}_{B_1} \left[\sqrt{\frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)}} \right] \\ &= \exp(-y/2) \int_{\mathbf{X}^k} \sqrt{\frac{\bar{g}_{B_2}(\mathbf{x}^k)}{\bar{g}_{B_1}(\mathbf{x}^k)}} g_{B_1}(\mathbf{x}^k) d\lambda^{\otimes kn}(\mathbf{x}^k). \end{aligned}$$

Now, by Assumption 4 it follows that $g_B \leq \bar{g}_B$ almost everywhere. Thus, we have

$$\begin{aligned} \mathbb{P}_{B_1} \left[\log \frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)} \geq y \right] &\leq \exp(-y/2) \int_{\mathbf{X}^k} \sqrt{\bar{g}_{B_2}(\mathbf{x}^k)} \sqrt{\bar{g}_{B_1}(\mathbf{x}^k)} d\lambda^{\otimes kn}(\mathbf{x}^k) \\ &\leq \exp(-y/2) \rho(\bar{P}_{B_2}^{\otimes k}, \bar{P}_{B_1}^{\otimes k}), \end{aligned}$$

where we are interpreting the definition of the Hellinger affinity function $\rho(\cdot, \cdot)$ as a function of two bounded positive measures, not necessarily probability measures.

At this point, we can follow the same argument as in Lemma 2 in [35], page 477, where the Hellinger affinity of two members of the convex hull of sets of probability distributions is shown to be less than the product of the Hellinger affinity of the factors. In our particular case, the measures $\bar{P}_B^{\otimes k}$ are not probability distributions, nonetheless, the same disintegration argument holds. Thus, we obtain

$$\rho(\bar{P}_{B_2}^{\otimes k}, \bar{P}_{B_1}^{\otimes k}) \leq \prod_{t=1}^k \prod_{j=1}^n \rho(\bar{P}_{B_2}^j, \bar{P}_{B_1}^j),$$

where \bar{P}_B^j is the measure with Radon-Nikodym derivative $\bar{g}_B(x) = \frac{1}{\mu_0(B)} \int_B (p_\theta^j(x))^{[A^{k-t}]_{ij}} d\mu_0(\theta)$ with respect to λ .

In addition, by Jensen's inequality⁵, with $x^{[A^{k-t}]_{ij}}$ being a concave function and $1/\mu_0(B) \int_B d\mu_0 = 1$, we have that

$$\bar{g}_B(x) \leq \left(\frac{1}{\mu_0(B)} \int_B p_\theta^j(x) d\mu_0(\theta) \right)^{[A^{k-t}]_{ij}}.$$

thus,

$$\mathbb{P}_{B_1} \left[\log \frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)} \geq y \right] \leq \exp(-y/2) \prod_{t=1}^k \prod_{j=1}^n \rho(P_{B_1}^j, P_{B_2}^j)^{[A^{k-t}]_{ij}},$$

where P_B^j is the probability distribution associated with the density $\frac{1}{\mu_0(B)} \int_B p_\theta^j(x) d\mu_0(\theta)$.

Assumption 3 and the compactness of Θ guarantees that $\rho(P_{B_1}^j, P_{B_2}^j) > \alpha$ for some positive α , thus similarly as in Lemma 2, we have that

$$\begin{aligned} \mathbb{P}_{B_1} \left[\log \frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)} \geq y \right] &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \prod_{t=1}^k \prod_{j=1}^n \rho(P_{B_1}^j, P_{B_2}^j)^{1/n} \\ &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \exp \left(-\frac{k}{n} \sum_{j=1}^n h^2(P_{B_1}^j, P_{B_2}^j) \right). \end{aligned}$$

Finally, by using the metric defined for the n -Hellinger ball and the fact that for a metric $d(A, B)$ for two sets A and B $d(A, B) = \inf_{x \in A, y \in B} d(x, y)$ we have

$$\begin{aligned} \mathbb{P}_{B_1} \left[\log \frac{\bar{g}_{B_2}(\mathbf{X}^k)}{\bar{g}_{B_1}(\mathbf{X}^k)} \geq y \right] &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \exp \left(-k \left(\sqrt{\frac{1}{n} \sum_{j=1}^n h^2(P_{B_1}^j, P_{B_2}^j)} \right)^2 \right) \\ &\leq \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1 - \delta} \right) \exp \left(-k \left(\sqrt{\frac{1}{n} \sum_{i=1}^n h^2(P_{\theta_1}^j, P_{\theta_2}^j)} - r^1 - r^2 \right)^2 \right). \end{aligned}$$

□

Lemma 3 provides a concentration result for the logarithmic ratio between two weighted densities over a pair of subsets B_1 and B_2 . The terms involving the auxiliary variable y and the influence of the graph, via δ are the same as in Lemma 2. Moreover, the rate at which this bound decays exponentially is influenced now by the radius of the two disjoint Hellinger balls where B_1 and B_2 are contained respectively.

The bound provided in Lemma 3 is defined for the random variables \mathbf{X}^k having a distribution $\mathbf{P}_B^{\otimes k}$. Nonetheless, \mathbf{X}^k are distributed according to $\mathbf{P}^{\otimes k}$. Therefore, we introduce a lemma that relates the Hellinger affinity of distributions defined over subsets of Θ .

Lemma 4 *Let Assumptions 1, 3 and 4 hold. Consider $\mathbf{P}_B^{\otimes k}$ as the distribution with density g_B as defined in Eq. (13), for $B \subseteq \mathcal{B}_R$. Then $h(\mathbf{P}_B^{\otimes k}, \mathbf{P}^{\otimes k}) \leq \sqrt{nk}R$.*

Proof By Jensen's inequality we have that

$$\sqrt{g_B(\mathbf{x})} \geq \frac{1}{\mu_0(B)} \int_B \sqrt{\prod_{t=1}^k \prod_{j=1}^n p_\theta^j(x_t^j)} d\mu_0(\theta).$$

⁵ For a concave function ϕ and $\int_\Omega f(x) dx = 1$, it holds that $\int_\Omega \phi(g(x)) f(x) dx \leq \phi(\int_\Omega g(x) f(x))$.

Then, by definition of the Hellinger affinity, it follows that

$$\rho(\mathbf{P}_B^{\otimes k}, \mathbf{P}^{\otimes k}) \geq \int_{\mathcal{X}^{\otimes k}} \sqrt{\prod_{t=1}^k \prod_{j=1}^n p^j(x_t^j)} \left(\frac{1}{\mu_0(B)} \int_B \sqrt{\prod_{t=1}^k \prod_{j=1}^n p_\theta^j(x_t^j)} d\mu_0(\theta) \right) d\lambda^{\otimes nk}(\mathbf{x}).$$

By using the Fubini-Tonelli Theorem, we obtain

$$\begin{aligned} \rho(\mathbf{P}_B^{\otimes k}, \mathbf{P}^{\otimes k}) &\geq \frac{1}{\mu_0(B)} \int_B \int_{\mathcal{X}^{\otimes k}} \sqrt{\prod_{t=1}^k \prod_{j=1}^n p^j(x_t^j)} \sqrt{\prod_{t=1}^k \prod_{j=1}^n p_\theta^j(x_t^j)} d\lambda^{\otimes nk}(\mathbf{x}) d\mu_0(\theta) \\ &= \frac{1}{\mu_0(B)} \int_B \prod_{t=1}^k \prod_{j=1}^n \rho(P^j, P_\theta^j) d\mu_0(\theta) \\ &= \frac{1}{\mu_0(B)} \int_B \prod_{t=1}^k \prod_{j=1}^n (1 - h^2(P^j, P_\theta^j)) d\mu_0(\theta). \end{aligned}$$

Finally, by the Weierstrass product inequality it follows that

$$\begin{aligned} \rho(\mathbf{P}_B^{\otimes k}, \mathbf{P}^{\otimes k}) &\geq \frac{1}{\mu_0(B)} \int_B \left(1 - \sum_{t=1}^k \sum_{j=1}^n h^2(P^j, P_\theta^j) \right) d\mu_0(\theta) \\ &= \frac{1}{\mu_0(B)} \int_B \left(1 - n \frac{1}{n} \sum_{t=1}^k \sum_{j=1}^n h^2(P^j, P_\theta^j) \right) d\mu_0(\theta) \\ &\geq \frac{1}{\mu_0(B)} \int_B (1 - nkR^2) d\mu_0(\theta), \end{aligned}$$

where the last line follows by the fact that any density P_θ , inside the n -Hellinger ball defined in the statement of the lemma, is at most at a distance R to \mathbf{P} . \square

Finally, before presenting our main result for compact sets of hypotheses, we will state an assumption regarding the necessary mass all agents should have around the correct hypothesis θ^* in their initial beliefs.

Assumption 5 *The initial beliefs of all agents are equal. Moreover, they have the following property: for any constants $C \in (0, 1]$ and $r \in (0, 1]$ there exists a finite positive integer K , such that*

$$\mu_0\left(\mathcal{B}_{\frac{C}{\sqrt{k}}}\right) \geq \exp\left(-k \frac{r^2}{32}\right) \quad \text{for all } k \geq K.$$

Assumption 5 implies that the initial beliefs should have enough mass around the correct hypothesis θ^* when we consider balls of small radius. Particularly, as we take Hellinger balls of radius decreasing as $O(1/\sqrt{k})$, the corresponding initial beliefs should not decrease faster than $O(\exp(-k))$.

The assumption can almost always be satisfied by taking initial beliefs to be uniform. The reason is that, in any fixed dimension, the volume of a ball of radius $O(1/\sqrt{k})$ will usually scale as a polynomial in $1/\sqrt{k}$, whereas we only need to lower bound it by a decaying exponential in k . For concreteness, we show how this assumption is satisfied by an example.

Example: Consider a single agent, with a uniform initial, belief receiving observations from a standard Gaussian distribution, i.e. $X_k \sim \mathcal{N}(0, 1)$. The variance is known and the agent would like to estimate the mean. Thus the models are $P_\theta = \mathcal{N}(\theta, 1)$. Now, the Hellinger distance can be explicitly written as

$$h^2(P, P_\theta) = 1 - \exp\left(-\frac{1}{4}\theta^2\right).$$

Therefore, the Hellinger balls of radius $1/\sqrt{k}$ will correspond to euclidean balls in the parameter space of radius

$$2\sqrt{\log\left(\frac{1}{1-\frac{1}{k}}\right)}.$$

Uniform initial belief indicates that $\mu_0\left(\mathcal{B}_{\frac{1}{\sqrt{k}}}\right) = \Theta\left(\frac{1}{\sqrt{k}}\right)$, which can be made larger than $\exp(-k\frac{r^2}{32})$ for sufficiently large k .

We are ready now to state our main result regarding the concentration of beliefs around θ^* for compact sets of hypotheses.

Theorem 2 *Let Assumptions 1, 3, 4 and 5 hold, and let $\sigma \in (0, 1)$ be a given probability tolerance level. Moreover, for any $r \in (0, 1]$, let $\{R_k\}$ be a decreasing sequence such that for $k = 1, \dots, R_k \leq \min\left\{\frac{\sigma}{2\sqrt{2kn}}, \frac{r}{4}\right\}$. Then, the beliefs $\{\mu_k^i\}$, $i \in V$, generated by the update rule in Eq. (7) have the following property: with probability $1 - \sigma$,*

$$\mu_{k+1}^i(\mathcal{B}_r) \geq 1 - \chi \exp\left(-\frac{k}{16}r^2\right) \quad \text{for all } i \text{ and all } k \geq \max\{N, K\}$$

where

$$N = \inf\left\{t \geq 1 \mid \exp\left(\log\frac{1}{\alpha}\frac{4\log n}{1-\delta}\right) \sum_{l=1}^{L-1} K_l \exp\left(-\frac{t}{32}r_{l+1}^2\right) < \frac{\sigma}{2}\right\},$$

with K as defined in Assumption 5, $\chi = \sum_{l=1}^{L-1} \exp(-\frac{1}{16}r_{l+1}^2)$ and $\delta = 1 - \eta/n^2$, where η is the smallest positive element of the matrix A .

Proof Lets start by analyzing the evolution of the beliefs on a measurable set B with $\theta^* \in B$. From Eq. (7) we have that

$$\begin{aligned} \mu_k^i(B) &= \frac{\int_B \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)}{\int_{\Theta} \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)} \\ &\geq 1 - \frac{\int_{B^c} \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)}{\int_B \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)}. \end{aligned}$$

Now lets focus specifically on the case where B is a n -Hellinger ball of radius $r > 0$ with center at θ^* . In addition, since $R_k < r$, we get

$$\mu_k^i(\mathcal{B}_r) \geq 1 - \frac{\int_{\mathcal{B}_r^c} \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)}{\int_{\mathcal{B}_{R_k}} \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)}.$$

Our goal will be to use the concentration result in Lemma 3. Thus, we can multiply and divide by $\mu_0(\mathcal{B}_{R_k})$ to obtain

$$\mu_k^i(\mathcal{B}_r) \geq 1 - \frac{\int_{\mathcal{B}_r^c} \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta)}{\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k) \mu_0(\mathcal{B}_{R_k})}$$

Moreover, we use the covering of the set \mathcal{B}_r^c to obtain,

$$\begin{aligned} \mu_k^i(\mathcal{B}_r) &\geq 1 - \sum_{l=1}^{L-1} \sum_{m=1}^{K_l} \int_{\mathcal{F}_{l,m}} \prod_{t=1}^k \prod_{j=1}^n p_{\theta}^j(X_t^j)^{[A^{k-t}]_{ij}} d\mu_0(\theta) \Big/ \bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k) \mu_0(\mathcal{B}_{R_k}) \\ &\geq 1 - \sum_{l=1}^{L-1} \sum_{m=1}^{K_l} \bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k) \mu_0(\mathcal{F}_{l,m}) \Big/ \bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k) \mu_0(\mathcal{B}_{R_k}). \end{aligned} \tag{15}$$

The previous relation defines a ratio between two densities, i.e. $\bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k)/\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k)$, both for the wighted likelihood product of the observations, where the numerator is defined over to the set $\mathcal{F}_{l,m}$ and the denominator with respect to the set \mathcal{B}_{R_k} .

Lemma 3 provides a way to bound term $\bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k)/\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k)$ with high probability, thus

$$\begin{aligned} \mathbb{P}_{\mathcal{B}_{R_k}} \left(\left\{ \mathbf{X}^k \left| \sup_{l,m} \log \frac{\bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k)}{\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k)} \geq y \right. \right\} \right) &\leq \sum_{l=1}^{L-1} \sum_{m=1}^{K_l} \mathbb{P}_{\mathcal{B}_{R_k}} \left(\log \frac{\bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k)}{\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k)} \geq y \right) \\ &\leq \sum_{l=1}^{L-1} \sum_{m=1}^{K_l} \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1-\delta} \right) \exp \left(-k \left(\sqrt{\frac{1}{n} \sum_{j=1}^n h^2(p_m^j, p^j)} - \delta_l - R_k \right)^2 \right) \\ &\leq \sum_{l=1}^{L-1} \sum_{m=1}^{K_l} \exp(-y/2) \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1-\delta} \right) \exp \left(-k (r_{l+1} - \delta_l - R_k)^2 \right). \end{aligned}$$

where p_m^j is the density of at the point $\theta = m \in S_{\varepsilon_l}$, where S_{ε_l} is the maximal ε_l separated set of \mathcal{F}_l as in Definition 4.

Particularly, lets use the covering proposed in [8], where $\delta_l = r_{l+1}/2$. From this choice of covering, we have that

$$\begin{aligned} r_{l+1} - \delta_l - R_k &> r_{l+1} - r_{l+1}/2 - r_{l+1}/4 \\ &= r_{l+1}/4 \end{aligned}$$

where we have used the assumption that $R_k \leq r/4$ or equivalently $R_k \leq r_l/4$ for all $1 \leq l \leq L$.

Thus, we can set $y = -\frac{k}{16}r_{l+1}^2$ and it follows that

$$\mathbb{P}_{\mathcal{B}_{R_k}} \left(\left\{ \mathbf{X}^k \left| \sup_{l,m} \log \frac{\bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k)}{\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k)} \geq y \right. \right\} \right) \leq \exp \left(\log \frac{1}{\alpha} \frac{4 \log n}{1-\delta} \right) \sum_{l=1}^{L-1} K_l \exp \left(-\frac{k}{16}r_{l+1}^2 \right). \quad (16)$$

The probability measure in Eq. (16) is computed for \mathbf{X}^k distributed according to $\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}$. Nonetheless, \mathbf{X}^k is distributed according to the (slightly different) $\mathbf{P}^{\otimes k}$. Our next step is to relate these two measures.

First, we have that for any distribution $\mathbf{P}_\theta \in \mathcal{B}_{R_k}$, from the Definition 3 of the n -Hellinger ball, it holds that

$$\sqrt{\frac{1}{n} \sum_{j=1}^n h^2(P_\theta^j, P^j)} \leq R_k,$$

and we relate the total variation distance and the Hellinger affinity as in Lemma 1 in [34]; for any measurable set A it holds that

$$\sup_A \left(\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}(A) - \mathbf{P}^{\otimes k}(A) \right)^2 \leq 1 - \rho^2(\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}, \mathbf{P}^{\otimes k}),$$

and by definition of the Hellinger affinity we have that

$$\begin{aligned} \sup_A \left(\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}(A) - \mathbf{P}^{\otimes k}(A) \right)^2 &= 1 - (1 - h^2(\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}, \mathbf{P}^{\otimes k}))^2 \\ &\leq 2h^2(\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}, \mathbf{P}^{\otimes k}), \end{aligned}$$

where first we have used the relation that for any $x \in \mathbb{R}$, it holds that $1 - (1 - x^2)^2 < 2x^2$. Then, from Lemma 4 we have that

$$\sup_A \left(\mathbf{P}_{\mathcal{B}_{R_k}}^{\otimes k}(A) - \mathbf{P}^{\otimes k}(A) \right)^2 \leq 2knR_k^2.$$

Therefore, by considering the measurable subset $\Gamma^k = \left\{ \mathbf{X}^k \left| \sup_{l,m} \log \frac{\bar{g}_{\mathcal{F}_{l,m}}(\mathbf{X}^k)}{\bar{g}_{\mathcal{B}_{R_k}}(\mathbf{X}^k)} \geq -\frac{k}{16}r_{l+1}^2 \right. \right\}$, we have that

$$\mathbb{P}(\Gamma^k) < \mathbb{P}_{\mathcal{B}_{R_k}}(\Gamma^k) + \sqrt{2kn}R_k$$

$$\leq \exp\left(\log \frac{1}{\alpha} \frac{4 \log n}{1-\delta}\right) \sum_{l=1}^{L-1} K_l \exp\left(-\frac{k}{16} r_{l+1}^2\right) + \frac{\sigma}{2}.$$

Furthermore, we are interested in finding a large enough k such that the probability described in Eq. (16) is at most σ . Thus, we define

$$N \geq \inf \left\{ t \geq 1 \mid \exp\left(\log \frac{1}{\alpha} \frac{4 \log n}{1-\delta}\right) \sum_{l=1}^{L-1} K_l \exp\left(-\frac{t}{16} r_{l+1}^2\right) < \frac{\sigma}{2} \right\}.$$

Moreover, from Eq. (15) we obtain that with probability $1 - \sigma$ for all $k \geq N$,

$$\begin{aligned} \mu_k^i(\mathcal{B}_r) &\geq 1 - \sum_{l=1}^{L-1} \sum_{m=1}^{K_l} \exp\left(-\frac{k}{16} r_{l+1}^2\right) \frac{\mu_0(\mathcal{F}_{l,m})}{\mu_0(\mathcal{B}_{R_k})} \\ &= 1 - \sum_{l=1}^{L-1} \exp\left(-\frac{k}{16} r_{l+1}^2\right) \frac{\mu_0(\mathcal{F}_l)}{\mu_0(\mathcal{B}_{R_k})} \\ &\geq 1 - \frac{1}{\mu_0(\mathcal{B}_{R_k})} \sum_{l=1}^{L-1} \exp\left(-\frac{k}{16} r_{l+1}^2\right). \end{aligned}$$

Now, let's define $\chi = \sum_{l=1}^{L-1} \exp\left(-\frac{1}{16} r_{l+1}^2\right)$, then it follows that

$$\begin{aligned} \mu_k^i(\mathcal{B}_r) &\geq 1 - \frac{1}{\mu_0(\mathcal{B}_{R_k})} \sum_{l=1}^{L-1} \exp\left(-\frac{k}{16} r_{l+1}^2\right) \\ &= 1 - \frac{1}{\mu_0(\mathcal{B}_{R_k})} \sum_{l=1}^{L-1} \exp\left(-\frac{1}{16} r_{l+1}^2\right) \exp\left(-\frac{k-1}{16} r_{l+1}^2\right) \\ &\geq 1 - \frac{1}{\mu_0(\mathcal{B}_{R_k})} \chi \exp\left(-\frac{k-1}{16} r^2\right), \end{aligned}$$

where the last inequality follows from $r_l \geq r$ for all $L \leq l \leq 1$. Finally, by Assumption 5 we have that, for all $k \geq K$

$$\begin{aligned} \mu_k^i(\mathcal{B}_r) &\geq 1 - \chi \exp\left(-\frac{k-1}{16} r^2 + \frac{k-1}{32} r^2\right) \\ &= 1 - \chi \exp\left(-\frac{k-1}{32} r^2\right), \end{aligned}$$

or equivalently $\mu_{k+1}^i(\mathcal{B}_r) \geq 1 - \chi \exp\left(-\frac{k}{32} r^2\right)$. \square

Analogous to Theorem 1, Theorem 2 provides a probabilistic concentration result for the agents' beliefs around a Hellinger ball of radius r with center at θ^* for sufficiently large k .

6 Conclusions

We have proposed an algorithm for distributed learning with both countable and compact sets of hypotheses. Our algorithm may be viewed as a distributed version of Stochastic Mirror Descent applied to the problem of minimizing the sum of Kullback-Leibler divergences. Our results show non-asymptotic geometric convergence rates for the beliefs concentration around the true hypothesis.

It would be interesting to explore how variations on stochastic approximation algorithms will produce new non-Bayesian update rules for more general problems. Promising directions include acceleration results for proximal methods, other Bregman distances or constraints within the space of probability distributions.

Furthermore we have modeled interactions between agents as exchanges of local probability distributions (i.e., beliefs) between neighboring nodes in a graph. An interesting open question is to understand to what extent this can be reduced when agents transmit only an approximate summary of their beliefs. We anticipate that future work will additionally consider the effect of parametric approximations allowing nodes to communicate only a finite number of parameters coming from, say, Gaussian Mixture Models or Particle Filters.

References

1. Acemoglu D, Nedić A, Ozdaglar A (2008) Convergence of rule-of-thumb learning rules in social networks. In: Proceedings of the IEEE Conference on Decision and Control, pp 1714–1720
2. Acemoglu D, Dahleh MA, Lobel I, Ozdaglar A (2011) Bayesian learning in social networks. *The Review of Economic Studies* 78(4):1201–1236
3. Alanyali M, Venkatesh S, Savas O, Aeron S (2004) Distributed bayesian hypothesis testing in sensor networks. In: Proceedings of the American Control Conference, pp 5369–5374
4. Aumann RJ (1976) Agreeing to disagree. *The Annals of Statistics* 4(6):1236–1239
5. Barbarossa S, Sardellitti S, Di Lorenzo P (2013) Distributed detection and estimation in wireless sensor networks. preprint arXiv:13071448
6. Beal MJ (2003) Variational algorithms for approximate Bayesian inference. University of London United Kingdom
7. Beck A, Teboulle M (2003) Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters* 31(3):167–175
8. Birgé L (2015) About the non-asymptotic behaviour of bayes estimators. *Journal of Statistical Planning and Inference* 166:67–77
9. Borkar V, Varaiya PP (1982) Asymptotic agreement in distributed estimation. *IEEE Transactions on Automatic Control* 27(3):650–655
10. Cooke R (1990) Statistics in expert resolution: A theory of weights for combining expert opinion. In: Cooke R, Costantini D (eds) *Statistics in Science, Boston Studies in the Philosophy of Science*, vol 122, Springer Netherlands, pp 41–72
11. Dai B, He N, Dai H, Song L (2015) Scalable bayesian inference via particle mirror descent. preprint arXiv:150603101
12. Dai B, He N, Dai H, Song L (2016) Provable bayesian inference via particle mirror descent. In: Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, pp 985–994
13. Darmon G (1935) Sur les lois de probabilité estimation exhaustive. *CR Acad Sci Paris* 260(1265):85
14. DeGroot MH (1974) Reaching a consensus. *Journal of the American Statistical Association* 69(345):118–121
15. Fox CW, Roberts SJ (2012) A tutorial on variational bayesian inference. *Artificial intelligence review* 38(2):85–95
16. Gale D, Kariv S (2003) Bayesian learning in social networks. *Games and Economic Behavior* 45(2):329–346
17. Gelman A, Carlin JB, Stern HS, Rubin DB (2014) *Bayesian data analysis*, vol 2. Chapman & Hall/CRC Boca Raton, FL, USA
18. Genest C, Zidek JV, et al (1986) Combining probability distributions: A critique and an annotated bibliography. *Statistical Science* 1(1):114–135
19. Ghosal S (1997) A review of consistency and convergence of posterior distribution. In: *Varanashi Symposium in Bayesian Inference*, Banaras Hindu University
20. Ghosal S, Ghosh JK, Van Der Vaart AW (2000) Convergence rates of posterior distributions. *Annals of Statistics* pp 500–531
21. Ghosal S, Van Der Vaart A, et al (2007) Convergence rates of posterior distributions for noniid observations. *The Annals of Statistics* 35(1):192–223
22. Gilardoni GL, Clayton MK (1993) On reaching a consensus using degroot’s iterative pooling. *The Annals of Statistics* 21(1):391–401
23. Golub B, Jackson MO (2010) Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics* pp 112–149
24. Gubner JA (1993) Distributed estimation and quantization. *IEEE Transactions on Information Theory* 39(4):1456–1459
25. Hill TP, Dall’Aglia M (2012) Bayesian posteriors without bayes’ theorem. preprint arXiv:12030251
26. Jadbabaie A, Lin J, Morse AS (2003) Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control* 48(6):988–1001

27. Jadbabaie A, Molavi P, Sandroni A, Tahbaz-Salehi A (2012) Non-bayesian social learning. *Games and Economic Behavior* 76(1):210–225
28. Jadbabaie A, Molavi P, Tahbaz-Salehi A (2013) Information heterogeneity and the speed of learning in social networks. *Columbia Business School Research Paper* (13-28)
29. Juditsky A, Rigollet P, Tsybakov AB, et al (2008) Learning by mirror averaging. *The Annals of Statistics* 36(5):2183–2206
30. Koopman BO (1936) On distributions admitting a sufficient statistic. *Transactions of the American Mathematical society* 39(3):399–409
31. Lalitha A, Sarwate A, Javidi T (2014) Social learning and distributed hypothesis testing. In: *IEEE International Symposium on Information Theory*, pp 551–555
32. Lalitha A, Javidi T, Sarwate A (2015) Social learning and distributed hypothesis testing. preprint arXiv:1410.4307 arXiv:1410.4307
33. Lan G, Nemirovski A, Shapiro A (2012) Validation analysis of mirror descent stochastic approximation method. *Mathematical programming* 134(2):425–458
34. LeCam L (1973) Convergence of estimates under dimensionality restrictions. *The Annals of Statistics* pp 38–53
35. LeCam L (1986) *Asymptotic Methods in Statistical Decision Theory*. Springer-Verlag, New York
36. Li J, Li G, Wu Z, Wu C (2016) Stochastic mirror descent method for distributed multi-agent optimization. *Optimization Letters* pp 1–19
37. Mossel E, Tamuz O (2010) Efficient bayesian learning in social networks with gaussian estimators. arXiv preprint arXiv:10020747
38. Mossel E, Sly A, Tamuz O (2014) Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields* 158(1-2):127–157
39. Nedić A, Lee S (2014) On stochastic subgradient mirror-descent algorithm with weighted averaging. *SIAM Journal on Optimization* 24(1):84–107
40. Nedić A, Olshevsky A (2015) Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control* 60(3):601–615
41. Nedić A, Olshevsky A, Uribe CA (2015) Fast convergence rates for distributed non-bayesian learning. preprint arXiv:150805161 1508.05161
42. Nedić A, Olshevsky A, Uribe CA (2015) Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs. In: *Proceedings of the American Control Conference*, pp 5884–5889
43. Nedić A, Olshevsky A, Uribe CA (2016) Distributed gaussian learning over time-varying directed graphs. In: *2016 50th Asilomar Conference on Signals, Systems and Computers*, pp 1710–1714
44. Nedić A, Olshevsky A, Uribe CA (2016) Distributed learning with infinitely many hypotheses. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp 6321–6326
45. Nedić A, Olshevsky A, Uribe CA (2016) Network independent rates in distributed learning. In: *Proceedings of the American Control Conference*, pp 1072–1077
46. Nedić A, Olshevsky A, Uribe CA (2016) A tutorial on distributed (non-bayesian) learning: Problem, algorithms and results. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp 6795–6801
47. Olfati-Saber R, Franco E, Frazzoli E, Shamma JS (2006) Belief consensus and distributed hypothesis testing in sensor networks. In: *Networked Embedded Sensing and Control*, Springer, pp 169–182
48. Olshevsky A (2014) Linear time average consensus on fixed graphs and implications for decentralized optimization and multi-agent control. preprint arXiv:1411.4186 arXiv:1411.4186
49. Qipeng L, Aili F, Lin W, Xiaofan W (2011) Non-bayesian learning in social networks with time-varying weights. In: *30th Chinese Control Conference (CCC)*, pp 4768–4771
50. Qipeng L, Jiuhua Z, Xiaofan W (2015) Distributed detection via bayesian updates and consensus. In: *34th Chinese Control Conference (CCC)*, pp 6992–6997
51. Rabbat M (2015) Multi-agent mirror descent for decentralized stochastic optimization. In: *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2015 IEEE 6th International Workshop on, IEEE, pp 517–520
52. Rabbat M, Nowak R (2004) Decentralized source localization and tracking wireless sensor networks. In: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol 3, pp 921–924
53. Rabbat M, Nowak R, Bucklew J (2005) Robust decentralized source localization via averaging. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol 5, pp 1057–1060
54. Rahimian MA, Shahrampour S, Jadbabaie A (2015) Learning without recall by random walks on directed graphs. preprint arXiv:150904332

55. Rahnama Rad K, Tahbaz-Salehi A (2010) Distributed parameter estimation in networks. In: Proceedings of the IEEE Conference on Decision and Control, pp 5050–5055
56. Rivoirard V, Rousseau J, et al (2012) Posterior concentration rates for infinite dimensional exponential families. *Bayesian Analysis* 7(2):311–334
57. Schwartz L (1965) On bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 4(1):10–26
58. Shahrampour S, Jadbabaie A (2013) Exponentially fast parameter estimation in networks using distributed dual averaging. In: Proceedings of the IEEE Conference on Decision and Control, pp 6196–6201
59. Shahrampour S, Rahimian M, Jadbabaie A (2015) Switching to learn. In: Proceedings of the American Control Conference, pp 2918–2923
60. Shahrampour S, Rakhlin A, Jadbabaie A (2016) Distributed detection: Finite-time analysis and impact of network topology. *IEEE Transactions on Automatic Control* 61(11):3256–3268
61. Su L, Vaidya NH (2016) Asynchronous distributed hypothesis testing in the presence of crash failures. University of Illinois at Urbana-Champaign, Tech Rep
62. Sun SL, Deng ZL (2004) Multi-sensor optimal information fusion kalman filter. *Automatica* 40(6):1017–1023
63. Tsitsiklis JN, Athans M (1984) Convergence and asymptotic agreement in distributed decision problems. *IEEE Transactions on Automatic Control* 29(1):42–50
64. Viswanathan R, Varshney PK (1997) Distributed detection with multiple sensors i. fundamentals. *Proceedings of the IEEE* 85(1):54–63
65. Walker SG (2006) Bayesian inference via a minimization rule. *Sankhyā: The Indian Journal of Statistics (2003-2007)* 68(4):542–553
66. Wang C, Chazelle B (2016) Gaussian learning-without-recall in a dynamic social network. arXiv preprint arXiv:160905990
67. Zellner A (1988) Optimal information processing and bayes’s theorem. *The American Statistician* 42(4):278–280
68. Zhu Y, Song E, Zhou J, You Z (2005) Optimal dimensionality reduction of sensor data in multisensor estimation fusion. *IEEE Transactions on Signal Processing* 53(5):1631–1639